

1 **BACKGROUND OF THE INVENTION**

2 **1. Co-pending Applications Containing Related Material**

3 The following co-pending applications, which are filed the same day as this
4 application, are hereby incorporated by reference: U.S. Application No. __/____, entitled
5 "Selecting a Server to Service Client Requests," and U.S. Application No. __/____,
6 entitled "Routing Client Requests to Back-End Servers."
7

8 **2. The Field of the Invention**

9 The present invention relates to redirecting requests for content that is stored on a
10 server. More specifically, the present invention relates to methods and computer program
11 products for using a front-end server in transparently redirecting a request for content such
12 that a client system is unaware of the redirection.
13

14 **3. The Prior State of the Art**

15 At times, a server may need to redirect requests for content. Redirection may occur
16 because (i) the content has moved, either temporarily or permanently, to another server,
17 (ii) multiple representations of the content are available and the client should choose from
18 them, (iii) the request should be made through a proxy, and etc. In the prior art, servers send
19 redirect responses to the client system and the client system is responsible for reissuing the
20 request for content to the server identified in the redirect response. However, sending
21 redirect responses to a client may not provide any benefit in some systems.

22 For example, some systems include a proxy server that provides a single point of
23 access to content stored on back end servers. The back end server is a server in the ordinary
24 sense, storing content that clients may access. The proxy or front end server provides an

1 access point for the content that is stored on one or more back end servers. From the client's
2 perspective, it appears as if the front end server is the source for all of the content stored at
3 the back end servers. However, clients are not necessarily precluded from accessing back
4 end servers directly.

5 Making requests for content through front-end servers provides several benefits.
6 From time to time, content stored on one back-end server may be moved to another. If a
7 client has been accessing the moved content directly, the client must alter future requests to
8 reflect the new location of the content. Although redirect responses from the old back-end
9 server may help inform clients that the content has moved, redirection requires additional
10 processing by the client and leads to slower response times for the requested content. Also,
11 redirect responses from back-end servers may be provided on a temporary basis. Requests
12 occurring after a certain time period may indicate simply that the requested content is not
13 available.

14 Front-end servers also facilitate local caching at the client system. Frequently,
15 content, such as email content, is cached at the client system to improve performance. Only
16 changes made to the email content is exchanged between the client and back-end server. If
17 email content is moved from one back-end server to another, the local caching is invalidated
18 and the full email content must be exchanged. By using a front-end server, all content
19 appears to come from a single source, regardless of the particular back-end server that
20 actually stores the content. Any movement of content from one back-end server to another
21 is transparent to the client. As a result, local caching remains valid even after content has
22 been moved.

23 However, redirection, as taught in the prior art, introduces problems when practiced
24 in a front-end server system. First, the front-end server is supposed to make the storage of

SUMMARY OF THE INVENTION

These and other problems with the prior art are overcome by the present invention, which is directed toward transparently redirecting a request for content such that a client system is unaware of the redirection. A front-end server receives requests for content stored on one or more back-end servers. By making the requests through the front-end server, the client does not know and does not care which back-end server stores the content being requested. As far as the client is concerned, the front-end server is the source of the requested content.

Content may be stored on one or more back-end servers. Private content, such as email, calendars, tasks, etc., typically is stored at only one back-end server, whereas public content, like applications, discussion groups, etc., typically is stored on multiple servers. At times, a front-end server may request data from a back-end server that the back-end server does not store. For example, if private content is moved, the back-end server generates a redirect response that identifies the back-end server storing the moved content. Rather than pass the redirect response back to the client as is taught in the prior art, the front-end server reissues the content request to the back-end server identified in the redirect response. Upon receiving the requested content from the redirect back-end server, the front-end sends the content to the client system.

Public content is likely to be stored on multiple back-end servers. However, the prior art allows for only a single server to be identified in a redirect response. When the front-end server receives a request for content from a client, the request is modified to indicate that the front-end is making the request to the back-end server. The back-end server receiving the request for content examines the request to determine if the request includes the front-end indicator. If present, the back-end server identifies a list of back-end servers

1 that store the requested content. By receiving a list of back-end servers, the front-end server
2 is able to load balance requests for content among all the back-end servers that can satisfy
3 the request. If the front-end indicator is not present, the back-end server identifies a single
4 back-end server in the redirect response. In identifying only a single back-end server, the
5 redirect response will be compatible with prior art clients making direct requests to back-end
6 servers.

7 Because the front-end server processes redirections, rather than passing them to the
8 client as taught in the prior art, the front-end server preserves the benefits of having
9 transparent back-end server storage. Specifically, content can move from one back-end
10 server to another without impacting the client.

11 Additional features and advantages of the invention will be set forth in the
12 description which follows, and in part will be obvious from the description, or may be
13 learned by the practice of the invention. The features and advantages of the invention may
14 be realized and obtained by means of the instruments and combinations particularly pointed
15 out in the appended claims. These and other features of the present invention will become
16 more fully apparent from the following description and appended claims, or may be learned
17 by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 illustrates an exemplary system that provides a suitable operating environment for the present invention;

Figure 2 shows a client system accessing email content that has moved from one back-end server to another;

Figure 3 show a client system accessing content that is stored at various back-end servers through a front-end server; and

Figures 4A and 4B are flowcharts illustrating exemplary methods for redirecting requests for content stored at a back-end server according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Client requests for data may be redirected from one server to another. Some redirections occur because data has been either temporarily or permanently moved to a new server. In the prior art, redirections are sent to the client system that initiated the request and the client reissues the request to the server indicated in the redirection. Where a proxy requests data for a client, the proxy returns redirection responses to the client for processing. However, redirection according to the prior art proves to be inadequate for certain systems. For example, U.S. Application No. __/____, filed on the same day herewith, entitled "Routing Client Requests to Back-End Servers," which is hereby incorporated by reference, describes the use of a front-end server to access content stored on one or more back-end servers. The specific back-end server storing requested content is unknown and transparent to the client.

One benefit of accessing content through a front-end server as described in the afore-referenced application is that the client system is insulated from having to identify the specific server storing the data being requested. The client requests data from the front-end server and the front-end server determines which back-end server should satisfy the request. Use of the front-end server also insures that local caching on the client remains valid when content moves from one back-end server to another because it appears to the client that the front-end server is the source of content.

With all requests for content being directed to the front-end server, redirection presents a problem because it identifies specific back-end servers. For example, a redirect response may not provide any meaningful information. In some systems, clients may be allowed access to back-end servers only through a front-end server. This makes the redirection information of no use to the client because reissuing the request directly to the

Sub
A2 7

004007 "ET 26.03.00

WORKMAN, NYDEGGER & SEELEY
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE
SALT LAKE CITY, UTAH 84111

1 identified back-end server will result in the request being denied. Even if the back-end
2 server is directly accessible to the client, a direct request to the back-end server prevents the
3 client from receiving any of the benefits offered by the front-end server arrangement.

4 In a Microsoft Exchange ® system, redirections may occur for a variety of reasons.
5 For example, Exchange public folder trees may be replicated across multiple back-end
6 servers to provide greater accessibility. However, the extent of content replication (in
7 contrast to folder tree replication) is controlled on a folder-by-folder basis. As a result, a
8 given back-end server may include the folder hierarchy, but not the folder contents. A
9 folder is termed "ghosted" when the folder is present in the public folder tree of a back-end
10 server, but the contents are not replicated on the back-end server.

11 The front-end server may direct requests for content based on the public folder tree.
12 If a request is for the contents of a folder that is ghosted, the back-end server issues a
13 redirect response that identifies other back-end servers where the content is replicated.
14 According to the present invention, the front-end server receives the redirect response from
15 the back-end server and reissues the redirected request rather than passing the redirect
16 response back to the client. When the front-end server receives the requested content from a
17 back-end server identified in the redirect response, the content is sent to the client. From the
18 client's perspective, the content's source is the front-end server. The client has no indication
19 that the redirection occurred.

20 Because replicated data may be stored on more than one server, the back-end server
21 ghosting a folder is capable of identifying a list of back-end servers in a redirect response.
22 However, prior art clients are unable to receive a list of servers in a redirect response. As a
23 result, front-end servers add a front-end indicator to any requests they direct to back-end
24 servers. If a back-end server receives a request with the front-end indicator, the back-end

1 server knows that a list may be returned. Returning a list of servers allows the front-end
2 server to load balance requests across all back-end servers capable of servicing a particular
3 request. U.S. Application No. __/____, filed on the same day herewith, entitled
4 "Selecting a Server to Service Client Requests," which is hereby incorporated by reference,
5 describes how a particular back-end server may be identified to load balance requests.
6 Otherwise, if the front-end indicator is not present, the back-end server returns a single
7 back-end server in the redirect response. In one embodiment the front-end indicator is a
8 string added to the standard HTTP "User-Agent" header and the server list is returned in an
9 extended HTTP 305 Use Proxy redirect response. However, the present invention is not
10 limited to the use of any particular HTTP headers or responses.

11 In Exchange systems it is also possible that access to the mailbox storage of an email
12 client will result in a redirect response. The front-end server accesses a global catalog to
13 identify the specific back-end server that stores email for a particular user. However, there
14 may be a short delay between the time an email storage is moved and the time that the move
15 is reflected in the global catalog. A request for email content made to the old back-end
16 server will result in a redirect response. As described above, the front-end server will
17 process the redirect response and issue a request for the content to the back-end server
18 identified in the redirection. As a general rule, email storage is located on only one
19 back-end server at a time. Redirect responses for email content, therefore, typically will
20 identify a single server rather than a list. Like email content, other types of private content
21 (content that is specific to a single user), such as calendars, task data, etc., is stored
22 ordinarily at only one back-end server.

23 By processing redirect responses, front-end servers allow clients to make requests for
24 content without knowing or caring about which particular back-end server stores the

Cont.
Sub
A3

OFFICE OF THE CLERK

WORKMAN, NYDEGGER & SEELEY
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE
SALT LAKE CITY, UTAH 84111

1 requested content. The present invention extends to both methods and computer program
2 products for transparently redirecting a request for content such that a client system is
3 unaware of the redirection. The embodiments of the present invention may comprise a
4 special purpose or general purpose computer including various computer hardware, as
5 discussed in greater detail below.

6 Embodiments within the scope of the present invention also include
7 computer-readable media for carrying or having computer-executable instructions or data
8 structures stored thereon. Such computer-readable media can be any available media that
9 can be accessed by a general purpose or special purpose computer. By way of example, and
10 not limitation, such computer-readable media may comprise RAM, ROM, EEPROM,
11 CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage
12 devices, or any other medium which can be used to carry or store desired program code
13 means in the form of computer-executable instructions or data structures and which can be
14 accessed by a general purpose or special purpose computer. When information is
15 transferred or provided over a network or another communications connection (either
16 hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer
17 properly views the connection as a computer-readable medium. Thus, any such a
18 connection is properly termed a computer-readable medium. Combinations of the above
19 should also be included within the scope of computer-readable media. Computer-executable
20 instructions comprise, for example, instructions and data which cause a general purpose
21 computer, special purpose computer, or special purpose processing device to perform a
22 certain function or group of functions.

23 Figure 1 and the following discussion are intended to provide a brief, general
24 description of a suitable computing environment in which the invention may be

1 implemented. Although not required, the invention will be described in the general context
2 of computer-executable instructions, such as program modules, being executed by
3 computers in network environments. Generally, program modules include routines,
4 programs, objects, components, data structures, etc. that perform particular tasks or
5 implement particular abstract data types. Computer-executable instructions, associated data
6 structures, and program modules represent examples of the program code means for
7 executing steps of the methods disclosed herein. The particular sequence of such executable
8 instructions or associated data structures represent examples of corresponding acts for
9 implementing the functions described in such steps.

10 Those skilled in the art will appreciate that the invention may be practiced in
11 network computing environments with many types of computer system configurations,
12 including personal computers, hand-held devices, multi-processor systems,
13 microprocessor-based or programmable consumer electronics, network PCs, minicomputers,
14 mainframe computers, and the like. The invention may also be practiced in distributed
15 computing environments where tasks are performed by local and remote processing devices
16 that are linked (either by hardwired links, wireless links, or by a combination of hardwired
17 or wireless links) through a communications network. In a distributed computing
18 environment, program modules may be located in both local and remote memory storage
19 devices.

20 With reference to Figure 1, an exemplary system for implementing the invention
21 includes a general purpose computing device in the form of a conventional computer 20,
22 including a processing unit 21, a system memory 22, and a system bus 23 that couples
23 various system components including the system memory 22 to the processing unit 21. The
24 system bus 23 may be any of several types of bus structures including a memory bus or

1 memory controller, a peripheral bus, and a local bus using any of a variety of bus
2 architectures. The system memory includes read only memory (ROM) 24 and random
3 access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic
4 routines that help transfer information between elements within the computer 20, such as
5 during start-up, may be stored in ROM 24.

6 The computer 20 may also include a magnetic hard disk drive 27 for reading from
7 and writing to a magnetic hard disk 39, a magnetic disk drive 28 for reading from or writing
8 to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to
9 removable optical disk 31 such as a CD-ROM or other optical media. The magnetic hard
10 disk drive 27, magnetic disk drive 28, and optical disk drive 30 are connected to the system
11 bus 23 by a hard disk drive interface 32, a magnetic disk drive-interface 33, and an optical
12 drive interface 34, respectively. The drives and their associated computer-readable media
13 provide nonvolatile storage of computer-executable instructions, data structures, program
14 modules and other data for the computer 20. Although the exemplary environment
15 described herein employs a magnetic hard disk 39, a removable magnetic disk 29 and a
16 removable optical disk 31, other types of computer readable media for storing data can be
17 used, including magnetic cassettes, flash memory cards, digital video disks, Bernoulli
18 cartridges, RAMs, ROMs, and the like.

19 Program code means comprising one or more program modules may be stored on the
20 hard disk 39, magnetic disk 29, optical disk 31, ROM 24 or RAM 25, including an operating
21 system 35, one or more application programs 36, other program modules 37, and program
22 data 38. A user may enter commands and information into the computer 20 through
23 keyboard 40, pointing device 42, or other input devices (not shown), such as a microphone,
24 joy stick, game pad, satellite dish, scanner, or the like. These and other input devices are

1 often connected to the processing unit 21 through a serial port interface 46 coupled to
2 system bus 23. Alternatively, the input devices may be connected by other interfaces, such
3 as a parallel port, a game port or a universal serial bus (USB). A monitor 47 or another
4 display device is also connected to system bus 23 via an interface, such as video adapter 48.
5 In addition to the monitor, personal computers typically include other peripheral output
6 devices (not shown), such as speakers and printers.

7 The computer 20 may operate in a networked environment using logical connections
8 to one or more remote computers, such as remote computers 49a and 49b. Remote
9 computers 49a and 49b may each be another personal computer, a server, a router, a network
10 PC, a peer device or other common network node, and typically include many or all of the
11 elements described above relative to the computer 20, although only memory storage
12 devices 50a and 50b and their associated application programs 36a and 36b have been
13 illustrated in Figure 1. The logical connections depicted in Figure 1 include a local area
14 network (LAN) 51 and a wide area network (WAN) 52 that are presented here by way of
15 example and not limitation. Such networking environments are commonplace in office-
16 wide or enterprise-wide computer networks, intranets and the Internet.

17 When used in a LAN networking environment, the computer 20 is connected to the
18 local network 51 through a network interface or adapter 53. When used in a WAN
19 networking environment, the computer 20 may include a modem 54, a wireless link, or other
20 means for establishing communications over the wide area network 52, such as the Internet.
21 The modem 54, which may be internal or external, is connected to the system bus 23 via the
22 serial port interface 46. In a networked environment, program modules depicted relative to
23 the computer 20, or portions thereof, may be stored in the remote memory storage device. It
24

Sub
A4

1 will be appreciated that the network connections shown are exemplary and other means of
2 establishing communications over wide area network 52 may be used.

3 Figure 2 shows a client system accessing email content that has moved from one
4 back-end server to another. Email server A 250 includes mailbox storage A 252 containing
5 email client A new mailbox 254, whereas email server B 260 includes mailbox storage
6 B 262 containing email client A old mailbox 264. Email server C 270 includes mailbox
7 storage C 272. Email server A 250, email server B 260, and email server C 270 are
8 examples of back-end servers. Email client A new mailbox 254 and email client A old
9 mailbox 264 are examples of resources available through front-end server 240. Global
10 catalog server 230 provides front-end server 240 with current information regarding the
11 content stored at each of the available back-end servers. Additionally, the front-end server
12 240 is stateless and does not require, for example, a hard drive for storing program data.
13 U.S. Application No. __/____, filed on the same day herewith, and entitled "Selecting a
14 Server to Service Client Requests," which is hereby incorporated by reference, provides
15 greater detail with respect to how a particular back-end server may be identified and how a
16 front-end server operates.

17 Email client A 210 uses front-end server 240 in accessing email content. Various
18 protocols are available that allow front-end server 240 to route content requests on behalf of
19 email client A 210. Internet message access protocol 4 ("IMAP4"), post office protocol
20 version 3 ("POP3"), and hypertext transfer protocol ("HTTP") are all protocols that
21 front-end server 240 can direct to a particular back-end server on behalf of email client
22 A 210. The present invention, however, is not limited to any particular protocol. It is only
23 necessary that the protocol allow for requests to be routed by a front-end server. For
24

1 example, messaging application programming interface ("MAPI") is a protocol that
2 currently will not allow front-end server 240 to route requests for email client A 210.

3 Comparing requests for content by email client A 210 with and without front-end
4 server 240 may be helpful in understanding the present invention. For simplicity, assume
5 that all requests use HTTP as their email protocol. Without front-end server 240, a request
6 by email client A 210 for the contents of the email client A old mailbox 264 might be
7 entered as "http://serverB/exchange/clientA" and a request by email client A 210 for the
8 contents of the email client A new mailbox 254 might be entered as
9 "http://serverA/exchange/clientA". Note that in each case, it was necessary to specify the
10 server where the mailbox contents were stored. In contrast, if the front-end server 240 were
11 named "mail" a request by email client A 210 for the contents of the email client A new
12 mailbox 254 might be entered as "http://mail/exchange/clientA".

13 By using front-end server 240, email client A 210 does not know and does not care
14 what back-end server stores the requested email content. When the email content is moved
15 from email client A old mailbox 264 on server B 260 to email client A new mailbox 254 on
16 server A 250, front-end server 240 will direct the request to the appropriate back-end server
17 automatically. Nevertheless, there may be a short time delay between when email content is
18 moved from server B 260 to server A 250 and when global catalog server 230 reflects the
19 move. During this time delay, front-end server 240 continues to direct requests for the email
20 content of email client A 210 to server B 260. Upon receiving a request, server B 260
21 generates a redirect response that indicates that the email content for email client A 210 now
22 is located at server A 250.

23 In contrast to the prior art, when front-end server 240 receives the redirect response,
24 front-end server 240 reissues the request for email content to server A 250 rather than

Sub
A5

1 passing the redirect response back to email client A 210. As shown in Figure 2, email client
2 A 210 may only access email content through front-end server 240. A redirect response
3 indicating that server A 250 stores the email content from email client A 210 would be of no
4 benefit to email client A 210 because server A 250 cannot be reached directly. Furthermore,
5 such a redirect response may confuse email client A 210 because no request for content was
6 made to server B 260, and therefore a redirect response from server B 260 is unexpected.
7 Front-end server 240 overcomes these problems with prior art redirection by processing the
8 redirect response without any indication to email client A 210 that any redirection has
9 occurred. U.S. Application No. __/____, filed on the same day herewith, entitled
10 "Routing Client Requests to Back-End Servers," which is hereby incorporated by reference,
11 provides further details on the operation of front-end server 240 in requesting email content
12 for email client A 210. The present invention should not be interpreted as requiring that
13 back-end server only be accessible through a front-end server.

14 The present invention is not limited to any particular protocol and it is not limited to
15 any particular content. Figure 3 shows client system 310 using front-end server 340 to
16 access content that is stored at various back-end servers. Back-end server 350 includes web
17 store 352 containing App A 382 and Folder A 392 that stores Data A 394. Back-end server
18 360 includes web store 362 containing App B 384 and Folder A 392 that stores Data A 394.
19 Back-end server 370 includes web store 372 containing App A 382 and Folder A 392.
20 However, Folder A 392 is dashed in web store 372 to indicate that the storage hierarchy is
21 present at back-end server 370, but the data (e.g., Data A 394) is not.

22 App A 382, App B 384, and Folder A 392 are examples of public content stored at a
23 back-end server. App A 382 is available in web stores 352 and 372, whereas App B 384 is
24 only available in web store 362. Folder A 392 represents a storage hierarchy for organizing

Cont
Sub
A5

DOCKET # 2016-00000000

WORKMAN, NYDEGGER & SEELEY
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE
SALT LAKE CITY, UTAH 84111

1 public content. Microsoft Exchange ® products allow the determination of whether data
2 corresponding to a folder will be stored on a particular server to be made on a
3 folder-by-folder basis. Folder A 392 in web store 372 is set not to store any data. One
4 reason for not storing the data for a folder is that the folder may be of minor significance for
5 the users of a particular server. Because the folder is seldom if ever accessed, it is better to
6 conserve the storage space that would otherwise be required for storing the folder's contents.

7 Nevertheless, by including the hierarchy on each server, it is possible to access the
8 contents of each folder on any server. A request for the contents of Folder A 392 made to
9 web store 372 will result in a redirect response that indicates where the contents of Folder
10 A 392 may be accessed. Most likely, the access will be slower than would be provided by
11 access to web store 372 because lower bandwidth communication links may be used, but
12 this added burden is a minor imposition given the relative infrequency of access to Folder
13 A 392 at back-end server 370. A folder within the storage hierarchy that does not include
14 the folder's contents is said to be "ghosted."

15 When client system 310 requests App A 382 from front-end server 340, global
16 catalog server 330 identifies a list of servers that includes back-end server 350 and back-end
17 server 370. Front-end server 340 uses an authentication token associated with authentication
18 credentials received from client system 310 as a key for performing a hashing operation over
19 the list of servers. The hashing operation insures that requests for content are evenly
20 distributed over the servers where the content is available and that a given request and
21 authentication token will identify the same server each time the request is made (as long as
22 the available servers do not change). By requesting data from the same back-end server
23 each time a given request is made, any user-specific state information may be stored and
24 recalled at the back-end server, such as a read/unread state for messages in a public folder.

1 Authentication tokens and authentication credentials may be stored in global catalog server
2 330 and/or cached in front-end server 340.

3 Client system 310 requesting App B 384 from front-end server 340, identifies a
4 single server, back-end server 360, from global catalog server 330. Like the mailbox
5 examples from Figure 2, no hashing operation occurs because the content only may be
6 requested from one back-end server. Thus, for both public and private content, it is possible
7 that the hashing operation will not be necessary to identify a particular back-end server
8 where the content request is directed.

9 Global catalog server 330 identifies back-end server 350, back-end server 360, and
10 back-end server 370, in response to a request for the contents of Folder A 392 from
11 front-end server 340. If the hashing operation identifies back-end server 350 or back-end
12 server 360, the request is directed to the identified back-end server as described above. But,
13 if back-end server 370 is identified, front-end server 340 must perform additional work in
14 order to direct the request to a back-end server that can satisfy the request because the
15 requested content is ghosted. When front-end server 340 directs the request to back-end
16 server 370, back-end server 370 generates a redirect response. The contents of the redirect
17 response depend on whether or not front-end 340 has modified the request to indicate that
18 the response is from a front-end server. If so, the redirect response includes a list of
19 back-end servers that can satisfy the request for the ghosted folder. Otherwise, the redirect
20 response includes a single back-end server capable of satisfying the request. U.S.
21 Application No. __/_____, filed on the same day herewith, and entitled "Selecting a Server
22 to Service Client Requests," which is hereby incorporated by reference, provides additional
23 information on identifying a back-end server when public content is ghosted on a back-end
24 server.

1 Figures 4A and 4B are flowcharts illustrating exemplary methods for redirecting
2 requests for content stored at a back-end server according to the present invention. The
3 dashed line between blocks 442 and 452 of Figure 4A indicates that the acts and steps above
4 the line are performed at a front-end server whereas the acts and steps below the line are
5 performed at a back-end server. A similar dashed line is included in Figure 4B. Returning
6 to Figure 4A, block 412 represents the act of the front-end server receiving a request for
7 content. As part of a step for authenticating the client (420), the front-end server may
8 perform the acts of requesting client authentication credentials (422) and receiving the
9 proper client authentication credentials (424) back from the client.

10 The front-end server adds a front-end indicator to the request (432) so that a
11 back-end server receiving the request will know that it was sent by a front-end server rather
12 than by a client. The step for querying (440A and 440B of Figure 4B) a back-end server for
13 the content may include the acts of directing the request to a particular back-end server (442)
14 and receiving a redirect response (444 of Figure 4B) from the particular back-end server. In
15 generating the redirect response, the back-end server performs the following acts. The
16 back-end server receives the request for content (452) and, in decision block 454, examines
17 the request for a front-end indicator. If no front-end indicator is present, the redirect
18 response includes a single redirect back-end server (462). With a front-end indicator
19 present, decision block 456 determines if multiple back-end servers can satisfy the request
20 for content. Here too, if only one back-end server is identified, the redirect response
21 includes a single redirect back-end server (462). However, if multiple back-end servers are
22 able to satisfy the content request, the redirect response includes a list of redirect back-end
23 servers (458). After the one or more redirect back-end servers have been identified, the
24 redirect response is sent to the front-end server (464 of Figure 4B).

1 As described earlier, as part of a step for querying the back-end server for requested
2 content (440A of Figure 4A and 440B), the front-end performs the act of receiving the
3 redirect response from the back-end server. In decision block 466, the front-end server
4 determines if the redirect response includes a list of back-end servers or a particular
5 back-end server. If a list was received, the front-end server receives an authentication token
6 (472) and hashes the list (474) with the authentication token as a key to identify a particular
7 back-end server. The acts of receiving an authentication token (472) and hashing a list of
8 back-end servers (474) may be included within a step for distributing a request for content to
9 one of the back-end servers in the list. The authentication token may be identified during
10 the step for authenticating the client (420 of Figure 4A).

11 Once a particular back-end server is identified, the front-end server redirects the
12 request to the particular back-end server (482) and receives the requested content from the
13 redirect back-end server (484). The acts of redirecting the request (482) and receiving the
14 requested content (484) are acts that may occur in a step for retrieving the requested content
15 from a particular redirect back-end server. After the content is received, the front-end server
16 sends the requested content to the client system (492).

17 The present invention may be embodied in other specific forms without departing
18 from its spirit or essential characteristics. The described embodiments are to be considered
19 in all respects only as illustrative and not restrictive. The scope of the invention is,
20 therefore, indicated by the appended claims rather than by the foregoing description. All
21 changes which come within the meaning and range of equivalency of the claims are to be
22 embraced within their scope.

23 What is claimed and desired to be secured by United States Letters Patent is:
24